Noname manuscript No.
(will be inserted by the editor)

Solving Visual Madlibs with Multiple Cues

Tatiana Tommasi · Arun Mallya · Bryan Plummer · Svetlana Lazebnik · Alexander C. Berg · Tamara L. Berg

arXiv:1611.00393v1 [cs.CV] 1 Nov 2016

Received: date / Accepted: date

Abstract This paper presents an approach for answering fill-in-the-blank multiple choice questions from the Visual Madlibs dataset. Instead of generic and commonly used representations trained on the ImageNet dataset, our approach employs a combination of networks trained for specialized tasks such as scene recognition, person activity classification, and attribute prediction. We also present a method for localizing phrases from candidate answers in order to provide spatial support for feature extraction. We map each of these features, together with candidate answers, to a joint embedding space through normalized canonical correlation analysis (CCA). Finally, we solve an optimization problem to learn to combine CCA scores from multiple cues to select the best answer. Extensive experimental results show a significant improvement over the previous state of the art and confirm that answering questions from a wide range of types benefits from examining a variety of image cues and carefully choosing the spatial support of feature extraction.

T. Tommasi
Dept. of Computer Control and Management Engineering
University of Rome, La Sapienza, Italy
E-mail: tommasi@dis.uniroma1.it

A. Mallya
University of Illinois at Urbana Champaign, IL, USA

B. A. Plummer
University of Illinois at Urbana Champaign, IL, USA

S. Lazebnik
University of Illinois at Urbana Champaign, IL, USA

A. C. Berg
University of North Carolina at Chapel Hill, NC, USA

T. L. Berg
University of North Carolina at Chapel Hill, NC, USA

Keywords Visual Question Answering · Cue Integration · Region Phrase Correspondence · Computer Vision · Language

1 Introduction

Visual Question Answering (VQA) has recently gained considerable popularity, with multiple new datasets being released (Antol et al, 2015; Malinowski and Fritz, 2014; Ren et al, 2015a; Gao et al, 2015; Zhu et al, 2016; Yu et al, 2015). What makes this task both attractive and challenging is that the range of questions that can be asked about images is very open-ended, potentially requiring a large amount of world knowledge as well as the ability to reason. The information needed to answer a question might be easily localized in a part of the image or might require accessing specific knowledge beyond the depicted content. Consider, for instance, a simple question about the position of an object: the answer could involve the overall scene (*e.g.*, it is in the kitchen), other reference objects (*e.g.*, it is on the table), their appearance (*e.g.*, it is against the blue wall), details about people (*e.g.*, it is in the girl's hand), activities (*e.g.*, it is floating in water) or even understanding of time and causality (*e.g.*, it is falling and about to touch the ground).

To date, a number of diverse solutions for VQA have been proposed, as surveyed in Section 2. Most existing approaches use a single representation computed by a deep Convolutional Neural Network (CNN), e.g., VGG-Net (Simonyan and Zisserman, 2014), trained on the ImageNet dataset (Russakovsky et al, 2015). However, it is hard to believe that a generic feature trained on a limited number of object classes can have sufficiently broad coverage and fine-grained discrimination power

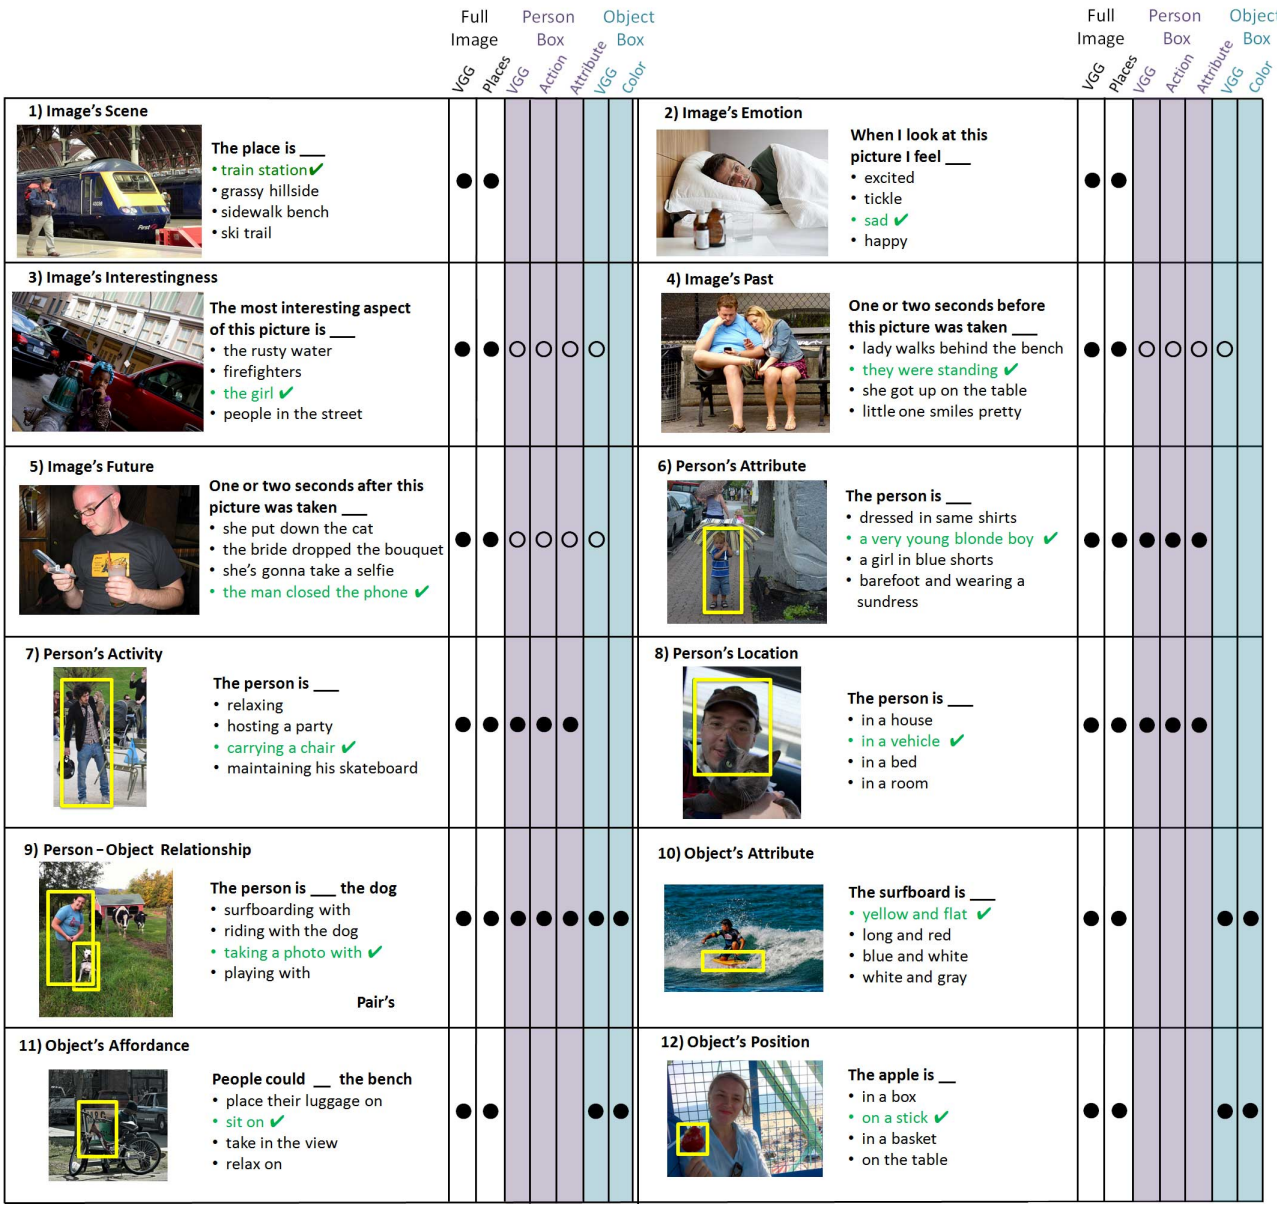
Full Image
Person Box
Object Box
VGG
Places
VGG
Action
Attribute
VGG
Color
1) Image's Scene
The place is ___
• train station ✔
• grassy hillside
• sidewalk bench
• ski trail
2) Image's Emotion
When I look at this picture I feel ___
• excited
• tickle
• sad ✔
• happy
3) Image's Interestingness
The most interesting aspect of this picture is ___
• the rusty water
• firefighters
• the girl ✔
• people in the street
4) Image's Past
One or two seconds before this picture was taken ___
• lady walks behind the bench
• they were standing ✔
• she got up on the table
• little one smiles pretty
5) Image's Future
One or two seconds after this picture was taken ___
• she put down the cat
• the bride dropped the bouquet
• she's gonna take a selfie
• the man closed the phone ✔
6) Person's Attribute
The person is ___
• dressed in same shirts
• a very young blonde boy ✔
• a girl in blue shorts
• barefoot and wearing a sundress
7) Person's Activity
The person is ___
• relaxing
• hosting a party
• carrying a chair ✔
• maintaining his skateboard
8) Person's Location
The person is ___
• in a house
• in a vehicle ✔
• in a bed
• in a room
9) Person – Object Relationship
The person is ___ the dog
• surfboarding with
• riding with the dog
• taking a photo with ✔
• playing with
Pair's
10) Object's Attribute
The surfboard is ___
• yellow and flat ✔
• long and red
• blue and white
• white and gray
11) Object's Affordance
People could __ the bench
• place their luggage on
• sit on ✔
• take in the view
• relax on
12) Object's Position
The apple is __
• in a box
• on a stick ✔
• in a basket
• on the table

Fig. 1 The Visual Madlibs dataset consists of 12 types of questions with fixed prompts, each concerned with the entire image (types 1-5), a specified person (types 6-9), or a specified object (types 9-12). For types 6-12, ground truth boxes of specified entities are provided as part of the question and are shown in yellow. Each question comes with four candidate answers, and only one (colored green, with a tick) is considered to be correct. To answer these varied questions, we use features computed on the whole image (VGG, Places), on person boxes (VGG, Action, Attribute) and on object boxes (VGG, Color). Details of the individual cues are given in Section 4. For each question type, circles mark the cues that are used by our final combination method. White circles indicate that the respective cues were computed on automatically selected person and object boxes, as no ground truth boxes were provided as part of the question. All the examples here come from the Hard question-answering setting (see Section 2).

needed to answer the wide variety of visual questions. To understand an image and answer questions about it, it is necessary to leverage a rich set of cues from different sources and to consider both global and local information. Accordingly, we propose in this paper a multi-cue approach for visual question answering and investigate its performance on fill-in-the-blank questions from the Visual Madlibs dataset (Yu et al, 2015).

The Visual Madlibs dataset consists of 12 different types of multiple choice questions, each with a fixed prompt, examples of which are shown in Figure 1. Of these question types, types 1-5 are based on high-level aspects of the whole image, namely predicting the scene, the emotion evoked by the image, likely past and future events, and the most interesting aspects of the image. Types 6-8 are based on characteristics of a spec-

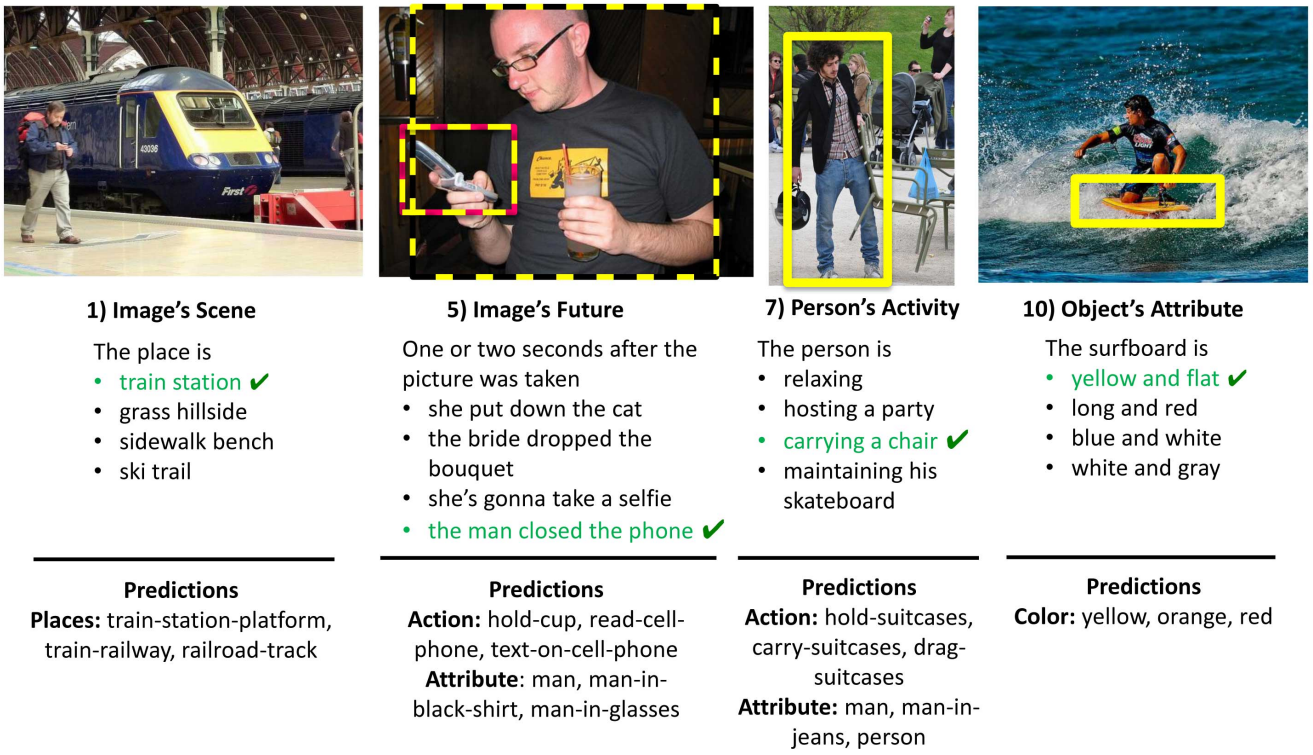

Fig. 2 Examples of four questions correctly answered by our system, along with intermediate predictions from our cue-specific deep networks. For each question, three top-scoring labels from the relevant networks are shown along the bottom. For the Future question, our method automatically selects the person and phone bounding boxes (shown with dashed lines), while for the Person's Activity and object's Attribute questions, bounding boxes are provided (solid yellow).

ified human subject, 9 is based on the interaction of a specified human and a specified object, while 10-12 are based on characteristics of a specified object. The person or object boxes that question types 6-12 focus on are provided as part of the question. Of the four possible answer options provided for each question, only one is considered to be correct (or most appropriate).

In order to answer the varied Visual Madlibs questions, we utilize a set of features extracted from image regions as well as from the whole image. These features come from "domain expert" models trained on specialized datasets. More specifically, we train a scene prediction network on the MIT Places dataset (Zhou et al, 2014); person action networks on the MPII (Pishchulin et al, 2014) and HICO (Chao et al, 2015) datasets; person attribute and object color networks on the Flickr30K Entities dataset (Plummer et al, 2015). Further details of these models will be presented in Section 4. In all, they give us the ability to predict a large vocabulary of semantically meaningful terms from the image, including close to 200 scene categories, 1000 actions and person-object interactions, 300 person attribute terms, and 11 colors.

Figure 1 shows which combination of cues is used for each question type, and Figure 2 gives a more detailed look at the cues computed by our expert models. It "zooms in" on four of the questions from Figure 1 and shows not only the question with the selected answer, but also the intermediate predictions of our scene, action, attribute, and color networks. The outputs of these networks are semantically interpretable and can help to understand why our system succeeds or fails on particular questions. Thus, in the Scene question example, we can see that the top scene label predictions from our Places network (*train-station-platform, train-railway, railroad-track*) are very similar to the correct answer (*train station*). For the Person's Activity question, our activity network cannot predict the correct activity (*carrying a chair*) even though it corresponds to an existing class; nevertheless, it is able to predict a sufficiently close class (*carry-suitcases*) to enable our image-text embedding to select the correct answer.

For our person- and object-centric networks, it is imperative to find the correct spatial support for feature extraction. While some Visual Madlibs question types come with pre-specified focus boxes, others do not – in particular, questions on Past, Future, and Interestingness. Consider the Interestingness question example in Figure 1 (question type 3). Two of the candidate answers for the most interesting aspect of this image are

the girl and *firefighters*. In order to score these answers, we need to attempt to localize the corresponding entities and determine whether they actually exist in the image. To this end, we utilize an automatic bounding box selection scheme that is described in Section 5. This scheme starts with candidate boxes produced by state-of-the-art person and object detectors (Liu et al, 2015; Ren et al, 2015b) and scores them using a region-phrase model trained on the Flickr30K Entities dataset (Plummer et al, 2015). The highest-scoring regions for phrases contained in answer choices then provide spatial support for feature extraction, and the region-phrase scores are also used as a component of the overall answer score.

To compute the compatibility between each of our network outputs and a candidate answer sentence or phrase, we use normalized Canonical Correlation Analysis (nCCA) (Gong et al, 2014), which maps the visual and textual features to a joint embedding space. In order to combine scores arising from different cues, we solve an optimization problem that learns the optimal combination weights. Details of our feature combination scheme will be given in Section 6. Section 7 will present our experimental results, which confirm that combining all our visual cues achieves the highest accuracy for most question types. In absolute terms, our results are state-of-the-art, outperforming the original Madlibs baseline (Yu et al, 2015), as well as a concurrent method (Mokarian et al, 2016) on all but one category.

A preliminary version of this work has appeared in BMVC (Tommasi et al, 2016). The journal version includes (1) a more detailed description and analysis of the different cues used for each question type, (2) a principled scheme for learning weights for combining multiple features, (3) extensive qualitative examples to better illustrate each part of the proposed approach, and (4) a statistical analysis of the coverage our cues provide for different types of Visual Madlibs questions (Section 7.1). The Visual Madlibs dataset project webpage has been updated with information on the features used in this work, as well as the validation set created for our experiments: `http://tamaraberg.com/visualmadlibs/`.

2 Related Work

Visual Question Answering. In Visual Question Answering (VQA), natural-language questions about an image are posed to a system, and the system is expected to reply with a short text answer. This task extends standard classification and image captioning, requiring techniques for multi-modal and knowledge-based reasoning for visual understanding. Initially proposed as a "Visual Turing Test" (Geman et al, 2015), the VQA format has been enthusiastically embraced as the basis for a number of specific datasets and benchmarks. The DAQUAR dataset (Malinowski and Fritz, 2014) is restricted to indoor scenes, while a number of more general datasets are based on MSCOCO images (Lin et al, 2014), including COCO-QA (Ren et al, 2015a), Baidu-FM-IQA (Gao et al, 2015), VQA (Antol et al, 2015), Visual7W (Zhu et al, 2016) and Visual Madlibs (Yu et al, 2015). Question-answer pairs can be generated automatically by NLP tools (Ren et al, 2015a), or created by human workers (Gao et al, 2015; Antol et al, 2015; Zhu et al, 2016; Yu et al, 2015). Assessing the quality of automatically generated free-form answers is not straightforward, thus most VQA benchmarks have a multiple-choice setting where performance can be easily measured as the percentage of correctly answered questions.

Among automatic methods for VQA, many combine CNNs and Long Short-Term Memory (LSTM) networks to encode the questions and output the answer (Gao et al, 2015; Malinowski et al, 2015; Andreas et al, 2016). Recent approaches also emphasize attention mechanisms for text-guided analysis of images (Zhu et al, 2016; Xu and Saenko, 2015; Yang et al, 2016; Shih et al, 2015, 2016). In contrast to these relatively complex methods, simple classification baselines (Jabri et al, 2016; Zhou et al, 2015) have also been shown to achieve high accuracy.

Fill-In-The-Blank Questions. Instead of asking explicit questions (*e.g.*, starting with who, what, where, when, why, which (Zhu et al, 2016)), we can ask systems to fill in incomplete phrases within declarative sentences. This is the strategy behind Visual Madlibs. As stated in the Introduction and shown in Figure 1, Visual Madlibs questions come in twelve distinct types, some with ground truth focus regions. The fact that each question has a well-defined type and structure that is known *a priori* makes the Visual Madlibs task more structured than general VQA, enabling us to reason up front about the types of features and processing needed to answer a given question. At the same time, due to the broad coverage and diversity of these question types, we can expect the types of cues that are useful for solving Visual Madlibs to also be useful for general VQA.

Visual Madlibs consists of 360,001 targeted natural language descriptions for 10,738 MSCOCO images, and fill-in-the-blank multiple choice questions are automatically derived from these descriptions. For each description type, the number of questions ranges between 4600 and 7500 and the descriptions contain more than three words on average. This makes Visual Madlibs notably different from VQA (Antol et al, 2015) and COCO-QA (Ren et al, 2015a) datasets, whose answers contain

around one word on average (see Zhu et al (2016), Table 1). An additional unique characteristic of Visual Madlibs is in the choice of the distractor (incorrect) answers, which have two levels of difficulty: in the Easy case the distractors are chosen randomly, while for the Hard case they are selected from the descriptions of images containing the same objects as the test image, with similar number of words as the correct answer, but not sharing with it any non-stop words.

Existing methods for answering Madlibs questions (Mallya and Lazebnik, 2016; Mokarian et al, 2016; Yu et al, 2015) have mainly used Canonical Correlation Analysis (CCA) (Hardoon et al, 2004; Hotelling, 1936) and normalized CCA (nCCA) (Gong et al, 2014) to create a multi-modal embedding where the compatibility of each putative answer with the image is evaluated. Mokarian et al (2016) have proposed CNN+LSTM models trained on Visual Madlibs, but these were not as accurate as CCA. Mokarian et al (2016) have also shown that the fill-in-the-blank task benefits from a rich image representation obtained by detecting several overlapping image regions, potentially containing different objects, and then average-pooling the CNN features extracted from them. This representation is able to cover the abundance of image details better than standard whole-image features, but it uses the same kind of descriptor at all image locations. In Section 7.3, we will demonstrate that our approach of using multiple specialized descriptors outperforms Mokarian et al (2016) for most question types.

Integrating External Knowledge Sources. Understanding images and answering visual questions often requires heterogeneous prior information that can range from common-sense to encyclopedic knowledge. To cover this need, some works integrate different knowledge sources either by leveraging training data with a rich set of different labels, or by exploiting textual or semantic resources such as DBpedia (Auer et al, 2007), ConceptNet (Liu and Singh, 2004) and WebChild (Tandon et al, 2014).

The approach adopted by Zhu et al (2015) learns a Markov Random Field model on scene categories, attributes, and affordance labels over images from the SUN database (Xiao et al, 2010). While this approach is quite powerful on the image side, the lack of natural language integration limits the set of possible questions that may be asked.

The method of Wu et al (2016) starts from multiple labels predicted from images and uses them to query Dbpedia. The obtained textual paragraphs are then coded as a feature through Doc2Vec (Le and Mikolov, 2014) and used to generate answers through an LSTM. A more sophisticated technique is proposed by Wang et al (2016) for an image question task that involves only answers about common-sense knowledge: the information extracted from images and knowledge-based resources is stored as a graph of inter-linked RDF triples (Lassila and Swick, 1999) and an LSTM is used to map the free-form text questions to queries that can be used to search the knowledge base. The answer is then provided directly as the result of this search, avoiding any limitations on the vocabulary that would otherwise be constrained by the words in the training set. Though quite interesting, both these approaches still rely on ImageNet-trained features, missing the variety of visual cues that can be obtained from networks tuned on tasks other than object classification.

As explained in the Introduction, our own approach to integrating external knowledge relies on training "expert" networks on specialized datasets for scenes, activities, and attributes. As one of the components of our approach, we use the CNN activity models developed in our ECCV 2016 paper (Mallya and Lazebnik, 2016), where we applied these models to Person Activity and Person-Object Relationship questions only.

3 Overview of Approach

To tackle multiple-choice fill-in-the-blank question answering we need a model that is able to evaluate the compatibility of each available answer choice $(a_1, \ldots, a_N)$ with the image and question pair (I, q). This necessitates a cross-modal similarity function that can produce a score $s(I, q, a)$ taking into consideration global (whole image to whole answer) and local (image region to phrase) correspondences, as well as multiple visual cues.

Our model has three main components: the image representation, the text representation, and a formulation for the cross-modal joint space and scoring function.

Representing the images. We introduce several feature types that depend on the question q and possibly on the specific answer choice a. This dependence is made explicit by choosing how to localize the feature extraction (*where* to compute the features) and *which features* to extract. Broadly speaking, we have the following four types of features:

- **Global image cues:** For all question types, we extract features from the whole image using our VGG ImageNet and Places networks (see Section 4 for details).
- **Cues from automatically selected boxes:** Question types 3-5 (Interestingness, Past, and Future) do not come with any ground truth person or

object boxes, but people and objects are often mentioned in candidate answers (see examples in Figure 1 and statistics in Section 7.1). We parse the candidate answers for mentioned entities and attempt to localize these using the procedure described in Section 5. Having found the best matching image region(s) for each mentioned entity, we extract specific features depending on the nature of the entity. In particular, for people, we extract bounding box VGG features as well as action and attribute features, and for objects, we extract bounding box VGG features only.

– **Cues from provided person boxes:** When dealing with person-centric questions (Types 6-9), we extract features from the person bounding box provided with the question. These include generic VGG features as well as features from our action and attribute networks.
– **Cues from provided object boxes:** For object-centric questions (Types 9-12) we extract features from the object bounding box provided with the question using our VGG and color networks.

Details of networks used for all of the above cues will be given in Section 4.

Representing the answers. Compared to our visual representation, our text representation is quite elementary. We employ the 300-dimensional word2vec embedding trained on the Google News dataset (Mikolov et al, 2013). Candidate answers are represented as the average of word2vec vectors over all the words. We represent out-of-vocabulary words using the null vector, and do not encode question prompts as they are identical for all questions of the same type (e.g., "the place is...").

Cross-modal embedding and scoring function. To learn a mapping from image and text features into a joint embedding space, we adopt standard normalized Canonical Correlation Analysis (nCCA) (Gong et al, 2014). For each question type, we obtain one or more nCCA scores for one or more cues corresponding to that type, and then form the final score as a linear combination of the individual scores with learned weights. Our cue combination and weight learning approaches are described in Section 6.

4 Cue-Specific Models

This section provides details of our cue-specific networks. For a complete summary of which networks are used for which question types, refer back to Figure 1.

Baseline network (VGG). For all question types, we use as our most basic global feature the output of the VGG-16 network (Simonyan and Zisserman, 2014) trained on 1000 ImageNet categories. We obtain a 4096-dimensional feature vector by averaging fc7 activations over 10 crops.

Scene network (Places). In addition to the above VGG feature, we have a global scene feature for each question type, derived from the Places VGG-16 network (Zhou et al, 2014). The MIT Places dataset contains about 2.5 million images belonging to 205 different scene categories. As with the baseline network, the Places network gives us 4096-dimensional fc7 features averaged over 10 crops.

Person activity networks (HICO, MPII). To describe person boxes for question types 3-9, we start by passing the boxes resized to 224×224 as input to the generic VGG-16 network. In order to obtain a more specialized and informative representation, we also use activity prediction networks trained on two of the largest currently available human activity image datasets: Humans Interacting with Common Objects (HICO) (Chao et al, 2015) and the MPII Human Pose Dataset (Pishchulin et al, 2014). HICO has 600 labels for different human-object interactions, *e.g.* ride-bicycle or repair-bicycle; the objects involved in the actions belong to the 80 annotated categories of the MSCOCO dataset (Lin et al, 2014). The MPII dataset has 393 categories, which include interactions with objects as well as solo human activities such as walking and running.

We employ the CNN architecture introduced in our previous work (Mallya and Lazebnik, 2016), which currently holds state-of-the-art classification accuracy on both the action datasets. This architecture is based on VGG-16 and it fuses information from a person bounding box and from the whole image. At training time it uses multiple instance learning to account for lack of per-person labels on the HICO dataset and a weighted loss to deal with unbalanced class distributions on both HICO and MPII. The model uses a weighted logistic loss in which mistakes on positive examples are weighted ten times more than the mistakes on negative examples, in order to offset the lack of balance in the dataset.

At test time, the network of Mallya and Lazebnik (2016) needs a person bounding box to provide a region of interest for feature extraction. For question types 6-9, these boxes are given in the ground truth. For question types 3-5, no boxes are given, so we use the automatic bounding box selection procedure that will be described in Section 5. In case of multiple people in an image, we run the network independently on each person and then average-pool the features. In case no person boxes

Act. : ride, stand-on-surfboard, surfing
Attr. : man, young-man, man-in-red-shirt

Act. : carry, hold-tennis-racket, hold-bat
Attr. : man, man-in-white-shirt, man-in-white-hat

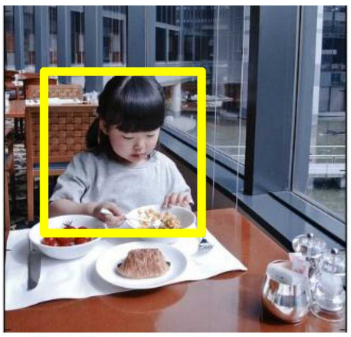

Act. : eat-at, sit-at-dining-table, lift-fork
Attr. : little-girl, young-girl, girl

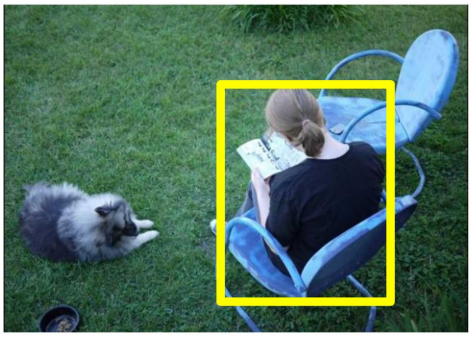

Act. : sit-on, lie-on-chair, read-book
Attr. : woman, girl, young-girl

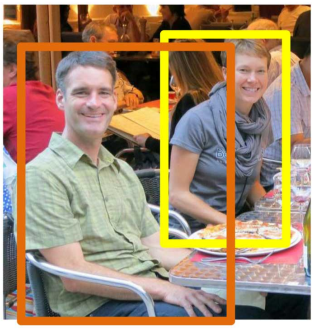

Act. Both Boxes: sit-at, eat-at-dining-table, hold-pizza
Attr. Box : man, young-man, guy
woman, girl, lady
Attr. Image: people, group-of-people, four-people

Act. Both Boxes: run, ride, straddle-horse
Attr. Box: jockey, man, rider
jockey, man, rider
Attr. Image: three-men, two-men, two people

Act. : ride, stand-on-surfboard, surfing
Attr. : man, young-man, man-in-red-shirt

Act. : brush-with-toothbrush, talking-on-cellphone, hold-cellphone
Attr. : little-boy, young-boy

Act. : wield, hold-baseball-bat, hold-tennis-racket
Attr. : baseball-player, man, boy

Fig. 3 Top three predicted person actions (Act.) and attributes (Attr.) for a few sample images. In the case of multiple people in an image, we specify the actions and attributes for specific boxes (underlined with the color of the box) as well as attributes for the whole image (Attr. Image). In the last row of images we show cases where action and attribute recognition fails.

are detected, we use the whole image as the region of interest.

Figure 3 presents some examples of class predictions of the activity networks. For various versions of our cue combination strategies, as described in Section 6, we will use either the fc7 activations of this network or the class prediction logits (inputs to the final sigmoid/softmax layer).

Person attribute network. For question types 3-9, alongside generic VGG features and activity features described above, we also extract high-level features based on a rich vocabulary of describable person attributes. To create such a vocabulary, we mine the Flickr30k Entities dataset (Plummer et al, 2015) for noun phrases that refer to people and occur at least 50 times in the training est. This results in

302 phrases that cover references to gender (man, woman), age (baby, elderly man), clothing (man in blue shirt, woman in black dress), appearance (Asian man, brunette woman), multiple people (two men, group of people), and more. An important advantage of our person attribute vocabulary is that it is an order of magnitude larger than those of other existing datasets (Sudowe et al, 2015; Bourdev et al, 2011). On the down side, attributes referring to males (*e.g.* man, boy, guy, *etc.*) occur twice as often as those referring to females (*e.g.* woman, girl, lady, *etc.*), and the overall class distribution is highly unbalanced (i.e., there are a few labels with many examples and many classes with just a few examples each).

We train a Fast-RCNN VGG-16 network (Girshick, 2015) to predict our 302 attribute labels based on person bounding boxes (in case of group attributes, the ground truth boxes contain multiple people). To compensate for unbalanced training samples, just as for the activity networks, we use a weighted logistic loss that penalizes mistakes on positive examples ten times more than on negative examples. Unlike our activity prediction network, our attribute network does not use global image context (we found that attribute predictions are much more highly localized and tend to be confused by outside context) and it predicts group attributes given a box with multiple people (such boxes naturally exist in the Flickr30K Entities annotations). As our labels are derived from natural language phrases, we manually grouped and ignored predictions on labels which could be simultaneously true but are not annotated in the dataset. For example, if a bounding box is referred to as *he*, *man in blue shirt*, *older man*, or *bald man*, related labels such as {*man, gentleman, guy, man in hard hat, asian man*} might also be true. Essentially, the presence of a label such as *he* does not conclusively indicate the absence of all other labels, such as *guy*, however it does indicate the absence of *she*, or *woman*. We manually created four such label groups representing man, woman, boy, and girl. If a label belongs to a given group, labels from all other groups can be safely considered as negatives, while labels within a group can be ignored while computing the training loss.

To give a quantitative idea of the accuracy of our person attribute prediction, the mAP of our network on the phrases of the Flickr30k test set that occur at least 50 (resp. 10) times is 21.98% (resp. 17.04%). We observe the following APs for some frequent phrases: man - 53.8%, woman - 51.3%, couple - 35.4%, crowd - 36.1%. It should be noted that these numbers likely underestimate the accuracy of our model. For one, they are based on exact matches and do not take synonyms into account. Moreover, there is a significant sparsity problem in the annotations, as numerous attribute phrases may be applicable to any person box but only a few are mentioned in captions. Qualitatively, the attribute labels output by our network are typically very appropriate, as can be seen from example predictions in Figure 3.

At test time, to obtain person bounding boxes from which to extract attribute features, we follow the same procedure as for the activity models described above. In case of multiple people boxes, the outputs of the attribute network are average-pooled. As with the activity models, either the inputs to the final sigmoid/softmax layer or the fc7 activations can be used for the downstream question answering task (refer to Sections 6 and 7 for details).

Color network. As described in Section 3, we extract object-specific cues for automatically detected boxes on question types 3-5 (Interestingness, Past, Future), as well as for provided focus boxes for question types 9-12. For all of those object boxes, just as for person boxes in question types 3-9, we extract generic VGG bounding box features. To complement these, we would also like to have a representation of object attributes analogous to our representation of person attributes. However, it is much harder to obtain training examples for a large vocabulary of predictable attributes for non-human entities. Therefore, we restrict ourselves to color, which is visually salient and frequently mentioned in Visual Madlibs descriptions, and is not captured well by networks trained for category-level recognition (Plummer et al, 2015). We follow Plummer et al (2015) and fine-tune a Fast-RCNN VGG-16 network to predict one of 11 colors that occur at least 1,000 times in the Flickr30K Entities training set: black, red, blue, white, green, yellow, brown, orange, pink, gray, purple. This network is trained with a one-vs-all softmax loss. The training is performed on non-person phrases to prevent confusion with color terms that refer to race. For our color feature representation we use the 4096-dimensional fc7 activation values extracted from the object bounding box.

Quantitative evaluation of a color network similar to ours can be found in Plummer et al (2015). The examples in Figure 4 provide a qualitative illustration of the color network outputs and indicate how color predictions may be helpful for answering object-attribute Visual Madlibs questions.

Note that we extract color features only from provided object boxes for questions 9-12. For questions 3-5, color is mentioned far more rarely in candidate answers; furthermore, automatically detected object boxes are much more noisy than person boxes making the color cues correspondingly unreliable.

Fig. 4 Examples of Object's Attribute questions with the top prediction of our Color network underneath. Even if the color mentioned in the answer is not among the ones predicted by our model, it can still be relevant (second row, first two images). The bottom right image is a failure case where the predicted color leads to the wrong answer.

5 Image Region Selection

Madlibs questions on Interestingness, Past, and Future do not provide a target image region. Consider the Future example in Figure 2, where each of the four candidate answers mentions a person and an object: *she put down the cat*, *the bride dropped the bouquet*, and so on. In order to pick the right choice, we need to select the best supporting regions for each of the entity mentions (*she*, *cat*, *bride*, *bouquet*) and use the respective matching scores as well as the features extracted from the selected regions as part of our overall image-to-answer scoring function.

We first parse all answers with the Stanford parser (Socher et al, 2013) and use pre-defined vocabularies to identify noun phrase (NP) chunks referring to a person or to an object. Then we apply the following region selection mechanisms for mentioned people and objects, respectively.

Person Box. We first detect people in an image using the Faster-RCNN detector (Ren et al, 2015b) with the default confidence threshold of 0.8. We discard all detected boxes with height or width less than 50 pixels since we find experimentally that these mainly contain noise and fragments. We also consider the smallest box containing all detected people, to account for cues originating from multiple people. Given the image and an answer, we attempt to select the box that best corresponds to the person mention in the answer. To this

end, we train a **Person CCA model** on the val+test set of Flickr30k Entities using person phrases (represented by average of word2vec) and person box features (302-dimensional vectors of predictions from our person attribute network of Section 4). As a lot of answer choices in the Madlibs dataset refer to people by pronouns or collective nouns such as *he, she, they, couple*, we augmented the training set by replacing person phrases as appropriate. For example, for phrases such as {*man, boy, guy, male, young boy, young man, little boy*}, we added training samples in which these phrases are replaced by *he* (and the same for *she*). Similarly, additional examples were created by replacing {*people, crowd, crowd of people, group of people, group of men, group of women, group of children*}, etc., with *they*, and {*two men, two women, two people*} with *couple*.

Given the trained Person CCA model, we compute the score for each person phrase from the candidate answer and each candidate person box from the image, and select the single highest-scoring box. A few example selections are shown in Figure 5. In case no words referring to people are found in a choice, all person boxes are selected.[1] The selected box provides spatial support for extracting person activity and attribute cues introduced in Section 4; in turn, these features, together with entire candidate answers (as opposed to just the person phrases), are used to train cue-specific CCA models as will be explained in the next section. The score of the Person CCA model for the selected box will also be used in a trained combination with the cue-specific CCA scores.

Object Box. We localize objects using the Single Shot MultiBox Detector (SSD) (Liu et al, 2015) that has been trained on the 80 MSCOCO object categories. SSD is currently the state of the art for detection in speed and accuracy. For each Visual Madlibs image, we consider the top 200 detections as object candidates and use the **Object CCA model** created for the phrase localization approach of (Plummer et al, 2015) to select the boxes corresponding to objects named in the sentences. This model is trained on the Flickr30k Entities dataset over Fast-RCNN fc7 features and average of word2vec features. We use the simplest model from that work, not including size or color terms. The top-scoring box from the image is used to extract object VGG features as will be explained in Section 6.

Figure 6 shows a few examples of object selection in action. As can be seen from the failure cases in the bottom row, object boxes selected by our method are less reliable than selected people boxes, since detection accuracies for general objects are much lower than for people and object boxes tend to be smaller. Therefore, instead of defining an object selection score based on the single highest-scoring region-phrase combination, as in the case of people above, we define a collective object score that will be used in the cue combination method of Section 6. Inspired by a kernel for matching sets of local features (Lyu, 2005), we take all of the $N = 200$ object boxes from the image and the M object phrases from the answer and then combine their CCA matching scores as follows:

$$K(image, answer) = \frac{1}{N}\frac{1}{M}\sum_{i=1}^{N}\sum_{j=1}^{M}\{\text{cos\_similarity}(box_i, phrase_j)\}^r \ , \quad (1)$$

where the parameter r assigns more relative weight to box-phrase pairs with higher similarity. We use $r = 5$ in our implementation.

6 Cue Combination

As described in Section 4, we extract several types of features from the images, aiming to capture multiple visual aspects relevant for different question types. How can we combine all these cues to obtain a single score $s(I, q, a)$ for each question and candidate answer?

The simplest combination technique is to concatenate 4096-dimensional fc7 features produced by each of our networks. In practice, due to the dimensionality of the resulting representation, we can only do this for a pair of networks, obtaining 8192-dimensional features. In our system, we mainly use this technique when we want to combine our baseline global VGG-16 network with one other cue.

To combine more than two features, we can stack lower-dimensional class prediction vectors (logits, or values before the final sigmoid/softmax layer). In particular, to characterize people, we concatenate the class predictions of HICO, MPII, and attribute networks, producing a compact feature vector of 1295 dimensions.

To enable even more complex cue integration, we learn CCA models on small subsets of cues and linearly combine their scores with learned weights. The following is a complete list of the individual CCA models used for our full ensemble approach:

– **Baseline + Places:** CCA trained on concatenated fc7 features from global ImageNet- and Places-trained VGG networks. This is used for all question types.

[1] Note that the images of the Visual Madlibs dataset are sampled from the MSCOCO dataset (Lin et al, 2014) to contain at least one person.

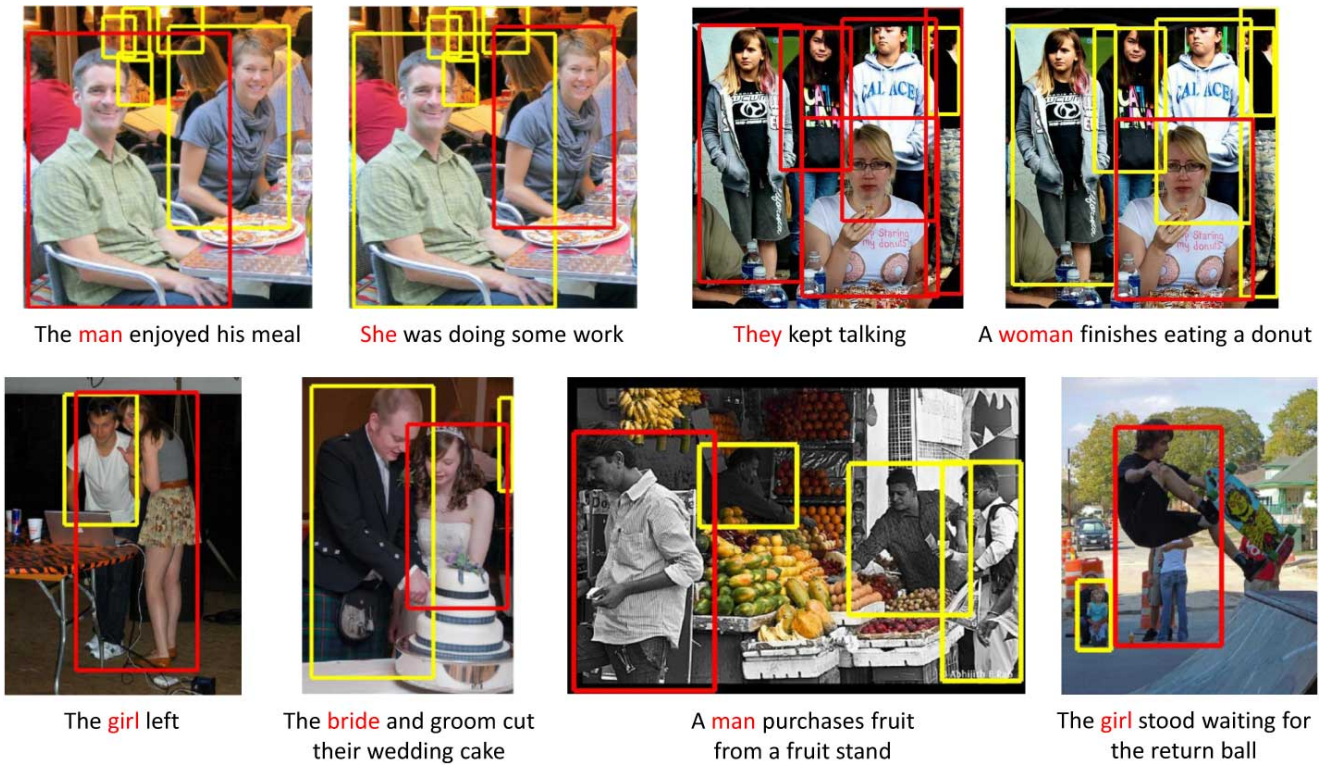
The man enjoyed his meal
She was doing some work
They kept talking
A woman finishes eating a donut
The girl left
The bride and groom cut their wedding cake
A man purchases fruit from a fruit stand
The girl stood waiting for the return ball

Fig. 5 Examples of selected person boxes based on person phrases. The person phrases are highlighted in red font and the corresponding selected boxes are also colored red. The yellow boxes are discarded either because they do not match the person mentioned in the phrase or because they are below the size threshold. In the third example from the left in the top row, CCA selects the overall box, thus all the person-specific boxes are colored red with the exception of the top right one which is discarded as it is below the size threshold. The last two images in the second row are failure cases.

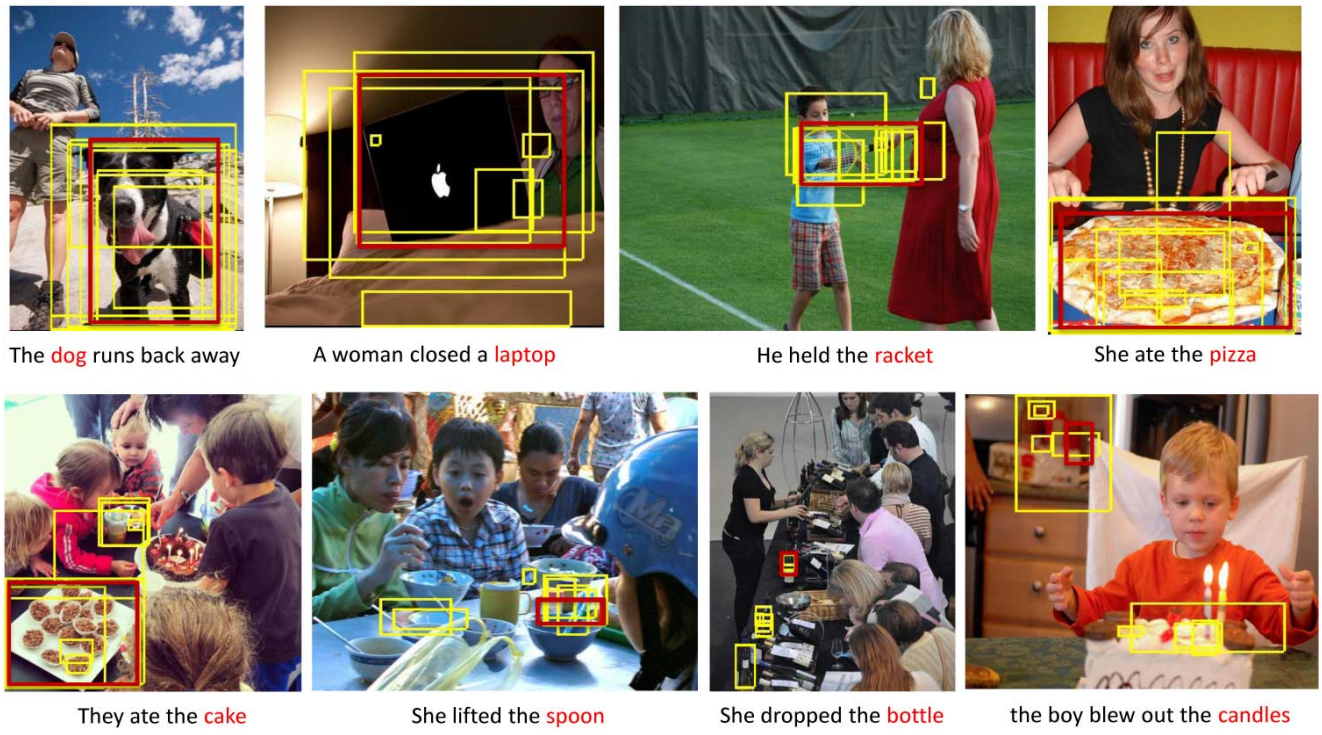
The dog runs back away
A woman closed a laptop
He held the racket
She ate the pizza
They ate the cake
She lifted the spoon
She dropped the bottle
the boy blew out the candles

Fig. 6 Examples of selected object boxes based on object phrases (in red). The red boxes are the top-scoring ones according to the object CCA model, while all the yellow boxes have lower scores. The top row presents correctly detected objects, while the bottom row shows failure cases.

- **Baseline + Person Box VGG:** CCA trained on concatenated fc7 features from VGG network applied to the whole image and person box. This cue is used for question types 3-5 (on automatically selected boxes) and 6-9 (on ground truth boxes). The reason for concatenating the global and person box features is to make sure that the resulting model is at least as strong as the baseline. The same reasoning applies to the other person-specific and object-specific models below.
- **HICO + MPII + Person Attribute:** CCA trained on concatenated logit scores from HICO, MPII, and Attribute networks. Used for question types 3-9.
- **Person selection score:** Person box selection score from the Person CCA model of Section 5. Used for question types 3-5.
- **Object selection score:** Scores from the Object CCA model of Section 5 combined using eq. (1). Used for question types 3-5.
- **Baseline + Object Box VGG:** CCA trained on concatenated fc7 features from the VGG network applied to the whole image and object box. Used for question types 3-5 (on automatically selected boxes) and 9-12 (on ground truth boxes).
- **Baseline + Object Box Color:** CCA trained on concatenated fc7 features from the VGG network applied to the whole image and color network applied to the object box. Used for question types 9-12.

To learn the combination weights, we divide the Visual Madlibs training set into an 80% training subset and a 20% validation subset. From the training subset, we learn the individual CCA models above using respective features and text descriptions[2]. For the validation set, we create three Easy and three Hard distractors for each correct description by following the same rules originally applied to create the test set (Yu et al, 2015).

For a particular question type, let s^j indicate the CCA score obtained on the validation sample (I, q, a) when using the jth model. We can then combine scores from all CCA models applicable to this question type as $S = \sum_j w^j s^j$. Let S_i denote the combined score for each candidate answer a_i for the considered sample, and i^* the index of the correct choice. We define the following convex loss:

$$L(S) = \max\{1 - S_{i^*} + \max_{i \neq i^*}\{S_i\}, 0\} \ . \quad (2)$$

This formulation assigns zero penalty when the score of the correct answer is larger by at least 1 than the scores of all the wrong choices. Otherwise, the loss is linearly proportional to the difference between the score of the correct answer and the maximum among the scores of the other choices. Over all the $k = 1, \ldots, K$ validation samples, we solve

$$\min_{\beta} \sum_{k=1}^{K} L(S)_k \quad \text{subject to} \quad \|w\|_1 \leq 1 \ , \quad w^j \geq 0 \ , \quad (3)$$

where the constraints specify that the weights for each feature should be positive, and the L1-norm condition can be seen as a form of regularization which induces a sparse solution. The optimization process can by easily implemented using the algorithm of Duchi et al (2008).

For a test question of a given type, we compute all the applicable CCA scores, combine them with the learned weights for that question type, and choose the answer with the highest combined score:

$$a_i^* = \operatorname*{argmax}_i\{S_i\} = \operatorname*{argmax}_i \left\{ \sum_j w^j s_i^j \right\} \ . \quad (4)$$

7 Experiments

In Section 7.1, to motivate our selection of cues for different questions, we examine the frequencies of cue-specific words in answers for each question type. In Section 7.2, we proceed to a detailed analysis of the multiple-choice answer task when using each cue separately. Finally, in Section 7.3 we evaluate the performance of the combined system.

7.1 Cue-Specific Category Statistics

Given the lists of 205 Places scene categories, 600 HICO action categories, 302 Flickr attribute categories, 80 MSCOCO object categories, and 11 color categories, we can compute the following statistics for correct answers (i.e., accurate descriptions) for each Visual Madlibs question type:

- **Madlibs coverage** = (number of answers that mention at least one of the categories) / (total number of answers);
- **Category list coverage** = (number of categories named at least once in the answers) / (total number of categories).

When counting the occurrences of the HICO actions, we consider past tense, continuous (-ing) and third person (-s) forms of the verbs. We also augment the MSCOCO object vocabulary with several word variants (e.g. bicycle, bike etc.) and singular/plural forms for all the objects.

[2] The Madlibs training set contains only the correct image descriptions, not the incorrect distractor choices.

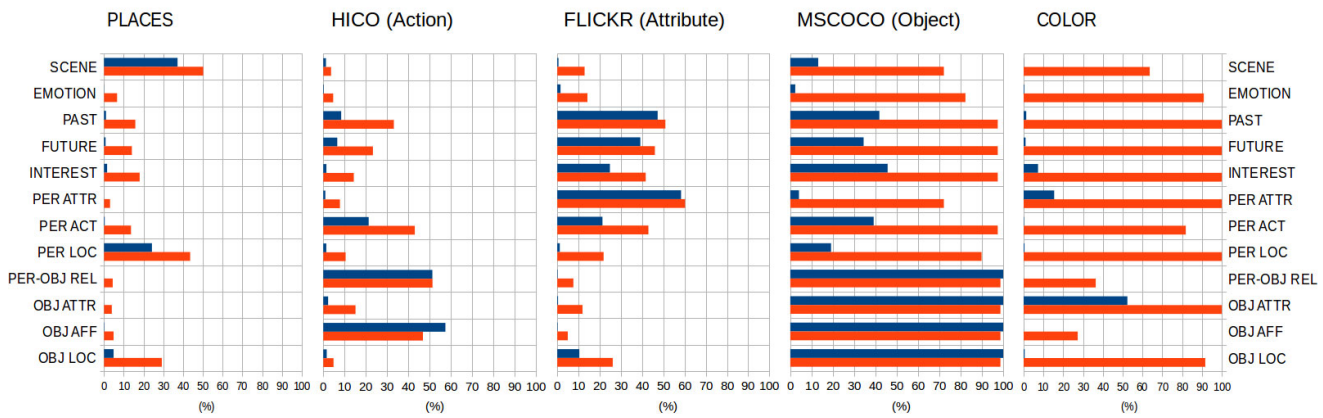

Fig. 7 Madlibs coverage (blue): indicates which percentage of the Visual Madlibs sentences mentions at least one of the Places (205 classes), action (HICO, 600 classes), attribute (302 classes), MSCOCO object (80 classes) and color (11 classes) categories. **Category list coverage (red)**: indicates which percentage of the category list is named at least once in the Visual Madlibs sentences.

Figure 7 shows the resulting statistics. Not surprisingly, Places categories have the best coverage on Scene questions: about 37% of the Visual Madlibs Scene answers mention one out of 50% of the Places categories. Beyond that, about 25% of Person Location answers mention one of 40% of the Places categories, and about 5% of Object Location answers mention one of 30% of these categories.

HICO action categories give the best coverage for Person-Object Relationship, Object Affordance, and Person Action questions. Flickr Attributes play important roles for Interestingness, Past, Future, Person Attribute, and Person Action questions. However, no more than about 50% (resp. 60%) of HICO Action (resp. Flickr Attribute) categories are mentioned in answers of any single given type.

By contrast, MSCOCO objects have 70-100% category list coverage for all question types, and 100% of the object-centric answers mention one of the MSCOCO categories. This is not surprising, since the Visual Madlibs dataset was created on top of MSCOCO images. However, despite this high coverage, we found that object-centric features played a limited role in our system outside of the object-centric question types (9-12), for which ground truth boxes are provided. For Scene and Emotion questions (1-2), object features are dominated by global cues, and for Person questions (6-9), they are dominated by Person features. As explained in Section 5, we do use object scores on automatically detected boxes for Interestingness, Past, and Future questions, but we will see in Section 7.3 that they provide only a minor improvement, most likely due to the high amount of uncertainty in detecting these objects.

Finally, Color categories play the most important role for Object Attribute questions: over 50% of answers for that question type mention a color, and 100% of the color names are mentioned. While a majority of the color names are also mentioned in all the other question types except for Person-Object Relationship and Object Affordance, the percentage of answers that actually mention a color is negligible. Our final system uses Color for all the object-centric questions (9-12), though as one would expect from the statistics, they only make a difference for the Object Attribute questions.

7.2 Single-Cue Results

This section analyzes the performance of our individual cues listed in Section 6. Table 1 presents the results on questions grouped by whole image (1-5), person-specific (6-9), and object-specific (10-12). The leftmost column shows the accuracy obtained with the baseline whole-image VGG fc7 feature. The subsequent columns show the performance obtained by concatenating this feature with the fc7 feature of each of our individual cue-specific network (as explained in Section 6, the reason for always combining individual cues with the baseline is to make sure they never get worse performance).

Whole-Image Questions. As shown in Table 1(a), using the Places features for Scene questions helps to improve performance over the VGG baseline. Emotion questions are rather difficult to answer but we can observe some improvement by adding Place features as well. We did not attempt to use person- or object-based features for the Scene and Emotion questions since the analysis of Section 7.1 indicated a negligible frequency

	Question Type		Full Image		Person Box				Object Box	
			Baseline VGG	B. + Places	B. + VGG	B. + HICO	B. + MPII	B. + Attr.	B. + VGG	B. + Color
(a)	1) Scene	Easy	87.73	**89.04**	–	–	–	–	–	–
		Hard	70.94	**73.22**	–	–	–	–	–	–
	2) Emotion	Easy	48.32	**49.53**	–	–	–	–	–	–
		Hard	35.50	**35.77**	–	–	–	–	–	–
	3) Interesting	Easy	78.11	78.74	**79.59**	79.47	78.55	79.31	78.86	–
		Hard	54.36	54.66	54.60	54.92	54.95	**56.02**	54.33	–
	4) Past	Easy	79.30	80.34	80.60	81.54	80.28	**81.68**	80.36	–
		Hard	53.89	53.95	55.37	55.09	54.11	**55.90**	53.78	–
	5) Future	Easy	79.52	80.10	80.76	**82.29**	80.42	81.30	80.61	–
		Hard	55.19	55.47	56.25	56.76	55.19	**57.58**	57.02	–
(b)	6) Person Attribute	Easy	53.32	54.10	59.00	54.10	55.60	**64.50**	–	–
		Hard	42.55	43.11	48.85	43.06	45.77	**54.64**	–	–
	7) Person Action	Easy	83.89	84.30	85.51	**87.46**	85.16	84.71	–	–
		Hard	67.56	68.10	69.47	**71.02**	70.03	68.68	–	–
	8) Person Location	Easy	84.59	**85.70**	84.50	85.29	84.51	84.33	–	–
		Hard	64.57	**66.66**	65.46	64.97	64.76	64.71	–	–
	9) Person Object Relationship	Easy	71.36	72.07	72.80	**75.77**	73.84	71.10	74.39	71.55
		Hard	54.46	54.65	56.84	**58.72**	56.84	54.48	55.85	54.58
(c)	10) Object Attribute	Easy	50.15	50.47	–	–	–	–	57.85	**59.50**
		Hard	44.99	45.62	–	–	–	–	53.63	**54.73**
	11) Object Affordance	Easy	80.56	82.76	–	–	–	–	**87.20**	82.32
		Hard	64.26	64.50	–	–	–	–	**67.65**	63.99
	12) Object Location	Easy	67.79	**69.41**	–	–	–	–	68.18	67.98
		Hard	56.46	**57.56**	–	–	–	–	57.34	56.43

Table 1 Accuracy on Madlibs questions with fc7 features. The Baseline VGG column gives performance for 4096-d fc7 outputs of the baseline network trained on ImageNet. For the columns labeled "B. + X", the baseline fc7 features are concatenated with fc7 features of different specialized networks, yielding 8192-d representations.

of person- and object-related words in the respective answers.

On the other hand, for Future, Past, and Interestingness questions, people and objects play an important role, hence we attempt to detect them in images as described in Section 5. From the selected person boxes we extract fc7 features from four different networks: the generic VGG network, the HICO and MPII Action networks, and the Person Attribute network trained on Flickr30K Entities. All of them give an improvement over the whole-image baseline, with the Attribute features showing the best performance in most cases. From the object regions we extract localized VGG features which also produce some improvement over the whole-image baseline in four out of six cases. Since, according to Figure 7, color is mentioned in only a tiny fraction of answers to the whole-image questions, we do not include it here.

Person Questions. For questions about specified people, Table 1(b) reports results with features extracted from the provided ground truth person box. Not surprisingly, Attribute features give the biggest improvement for Attribute questions, and HICO Action features give the biggest improvement for Action and Person-Object Relationship questions (recall that HICO classes correspond to interactions between people and MSCOCO objects). For the latter question type, the ground truth object region is also provided; by extracting the VGG and Color features from the object box we obtain accuracy lower than that of the HICO representation but still higher than that of the whole-image baseline. Finally, for Person Location questions, the global Places features work the best.

Object Questions. For questions about specified objects, Table 1(c) reports results with features extracted from the provided ground truth object box. We can see that Color features work best for Object Attribute questions, VGG features work best for Affordance questions, and Places features work best for Object Location questions.

7.3 Multi-Cue Results

Table 2 shows the results obtained by integrating multiple cues in a variety of ways. We exclude Scene and Emotion questions from the subsequent analysis: based on Figure 7, very few of their answers involve persons and objects, thus, our final cue combination for these question types is simply the concatenation of VGG and Places as shown in Table 1.

For ease of comparison, the first and second columns of Table 2 repeat the baseline and highest results from

	Question Type		fc7 Combination			Label Combination		CCA Score Combination		
			Baseline VGG	Baseline + Single Best Cue		HICO + MPII	HICO + MPII + Attr.	+ Person Score	+ Object Score	CCA Ensemble
(a)	3) Interesting	Easy	78.11	HICO	79.47	79.25	79.94	80.59	80.88	**82.34**
		Hard	54.36	Attr.	56.02	54.11	55.37	56.25	56.31	**57.92**
	4) Past	Easy	79.30	Attr.	81.68	82.17	84.09	84.17	84.97	**85.91**
		Hard	53.89	Attr.	55.90	55.23	58.17	58.29	59.60	**61.33**
	5) Future	Easy	79.52	HICO.	82.29	82.89	84.97	84.97	85.47	**86.63**
		Hard	55.19	Attr.	57.58	56.87	59.98	60.05	61.91	**62.73**
(b)	6) Person Attribute	Easy	53.32	Attr.	64.50	59.37	68.43	–	–	**68.68**
		Hard	42.55	Attr.	54.64	46.61	56.17	–	–	**56.38**
	7) Person Action	Easy	83.89	HICO	87.46	87.23	87.26	–	–	**88.43**
		Hard	67.56	HICO	71.02	71.35	71.42	–	–	**71.68**
	8) Person Location	Easy	84.59	Places	85.70	84.56	84.51	–	–	**86.28**
		Hard	64.57	Places	**66.66**	62.82	62.46	–	–	**66.66**
	9) Person Object Relationship	Easy	71.36	HICO	75.77	75.42	75.66	–	–	**77.08**
		Hard	54.46	HICO	**58.72**	56.68	56.88	–	–	57.92
(c)	10) Object Attribute	Easy	50.15	Color	59.50	–	–	–	–	**59.62**
		Hard	44.99	Color	**54.73**	–	–	–	–	**54.73**
	11) Object Affordance	Easy	80.56	Obj. VGG	87.20	–	–	–	–	**87.21**
		Hard	64.26	Obj. VGG	67.65	–	–	–	–	**67.69**
	12) Object Location	Easy	67.79	Places	69.41	–	–	–	–	**69.71**
		Hard	56.46	Places	57.34	–	–	–	–	**58.16**
	Average		64.34		68.66	–	–	–	–	**70.35**

Table 2 Results of combining multiple cues. Columns marked "fc7 Combination" give key results from Table 1 for reference. Columns marked "Label Combination" show results with combining the class activation vectors from the respective networks. Columns marked "+ Person Score" and "+ Obj. Score" show the results of a learned combination of the HICO + MPII + Attr. CCA with the region selection scores of Section 5. The CCA Ensemble column shows the results of combining all CCA scores appropriate for each question type (see Table 3 for the weights learned for this combination).

Table 1. The subsequent columns show performance obtained with other cue combinations.

The Label Combination columns of Table 2 show the results of concatenating the class prediction vectors from the HICO and MPII networks, and from all three person-centric networks (HICO+MPII+Attribute). For HICO+MPII, we observe a small drop in performance over the single best cue on whole-image questions (*i.e.*, in Interesting, Past, Future rows) and location-related questions (Person Location and Person-Object Relation), probably owing to the reduced feature dimension and loss of global contextual information as compared to the 8192-dimensional fc7 combination feature. On the other hand, HICO+MPII produces results comparable with the best fc7 cue for the Person Action question while being much more compact (993 vs. 8192 dimensions). By adding the attribute labels (HICO + MPII+Attribute column), we further improve performance, particularly on the Person Attribute question.

Recall from Section 5 that for Interestingness, Past, and Future questions, we perform focus region selection and compute Person and Object scores measuring the compatibility of person and object mentions in answers with the selected regions. These scores also provide some useful signal for choosing the correct answer, so we use the procedure of Section 6 to learn to combine each of them with the scores from the HICO+MPII+Attribute CCA model. For these two-cue problems, the learning procedure assigns a high weight to the combined action and attribute representation ($w^{HICO+MPII+Attribute} \geq 0.9$) and a small one to the Person and Object scores ($w^{Person/Obj.\ Selection} \leq 0.1$). The resulting accuracies are reported in columns labeled "+ Person Score" and "+ Object Score" of Table 2, and they show small but consistent accuracy improvements over the HICO+MPII+Attribute model, particularly for the hard questions.

The last column of Table 2 gives the performance of the full ensemble score using all the CCA models applicable to a given question type (refer back to Section 6 for the list of models). For completeness, Table 3 shows the learned per-cue weights for each question type. Generally, the most informative cues for each question type get assigned higher weights (*e.g.* HICO+MPII+Attribute features get high weights for Person Action and Attribute questions, but not for Person Location questions). From the last row of Table 2, we can observe an average improvement of 1.5% in accuracy for most of the questions with respect to the single best cue and 6% with respect to the baseline.

To date, the strongest competing system on Visual Madlibs is that of Mokarian et al (2016)[3]. Our CCA

[3] They use almost the same experimental protocol as us, except that they train their nCCA models on the entire Vi-

	Question Type		Full Image	Person Box			Object Box		
			B. + Places	B. + VGG	HICO + MPII + Attr.	Person Score	B. + VGG	B. + Color	Object Score
(a)	3) Interesting	Easy	0.00	0.00	0.64	0.00	0.36	–	0.00
		Hard	0.09	0.15	0.41	0.11	0.23	–	0.01
	4) Past	Easy	0.01	0.02	0.63	0.05	0.29	–	0.00
		Hard	0.04	0.05	0.48	0.16	0.17	–	0.09
	5) Future	Easy	0.03	0.01	0.68	0.04	0.24	–	0.00
		Hard	0.05	0.19	0.31	0.14	0.09	–	0.22
(b)	6) Person Attribute	Easy	0.00	0.21	0.79	–	–	–	–
		Hard	0.03	0.35	0.62	–	–	–	–
	7) Person Action	Easy	0.08	0.07	0.85	–	–	–	–
		Hard	0.00	0.48	0.52	–	–	–	–
	8) Person Location	Easy	0.67	0.00	0.34	–	–	–	–
		Hard	1.00	0.00	0.00	–	–	–	–
	9) Person Object Relationship	Easy	0.11	0.17	0.42	–	0.23	0.07	–
		Hard	0.03	0.39	0.13	–	0.44	0.01	–
(c)	10) Object Attribute	Easy	0.00	–	–	–	0.18	0.82	–
		Hard	0.00	–	–	–	0.20	0.80	–
	11) Object Affordance	Easy	0.20	–	–	–	0.80	0.00	–
		Hard	0.29	–	–	–	0.71	0.00	–
	12) Object Location	Easy	0.84	–	–	–	0.16	0.00	–
		Hard	0.77	–	–	–	0.23	0.00	–

Table 3 Weights assigned by the CCA score combination (Ensemble) method to each cue.

Ensemble results are superior to theirs on every question type except for Easy Object Attribute questions (we get 59.3% vs. their 62.4%) and Easy Object Location questions (69.2% vs. 77.4%). We attribute this difference to the use of a large set of object boxes in (Mokarian et al, 2016), while our approach focuses only on the the ground truth object position. This indicates that pooling features over multiple regions may be a good strategy even in cases when ground truth focus regions are provided.

Finally, Figure 8 shows answer choices selected with individual cues for the same questions that were originally shown in Figure 1, while Figure 9 shows a few failure cases.

8 Conclusions

We have shown that features representing different types of image content are helpful for answering multiple choice questions, confirming that external knowledge can be successfully transferred to the VQA task through the use of deep networks trained on specialized datasets. Further, through the use of an ensemble of CCA models, we have created a system that beats the previous state of the art on the Visual Madlibs dataset for most question types.

A detailed analysis of our approach has shown where further work would be beneficial. Person and object localization may be improved by a better interpretation of the sentences that does not focus only on separate entities, but understands their relationships and translates them into spatial constraints to guide region selection and feature extraction. Pooling features extracted from multiple object boxes in each image could also be beneficial. And, of course, training joint image-text models that can better deal with rare and unusual inputs remains an important open problem, as exemplified by the questions in the left column of Figure 9.

In the future, we are interested in extending our approach to more open-ended and general VQA tasks that do not rely on pre-specified question templates or provided ground truth focus regions.

Acknowledgments. This material is based upon work supported by the National Science Foundation under grants 1302438, 1563727, 1405822, 1444234, Xerox UAC, Microsoft Research Faculty Fellowship, and the Sloan Foundation Fellowship.

References

Andreas J, Rohrbach M, Darrell T, Klein D (2016) Deep compositional question answering with neural module networks. In: CVPR

Antol S, Agrawal A, Lu J, Mitchell M, Batra D, Zitnick CL, Parikh D (2015) VQA: Visual Question Answering. In: ICCV

sual Madlibs training set, without setting aside any validation data. We did the same in the original conference paper version of our work (Tommasi et al, 2016), where we did not perform weight learning and thus did not need the validation set.

Full Image
Person Box
Object Box
VGG
Places
VGG
Action
Attribute
VGG
Color
CCA Ensemble
1) Image's Scene
The place is ___
• train station ✔
• grassy hillside
• sidewalk bench
• ski trail
2) Image's Emotion
When I look at this picture I feel ___
• excited
• tickle
• sad ✔
• happy
3) Image's Interestingness
The most interesting aspect of this picture is ___
• the rusty water
• firefighters
• the girl ✔
• people in the street
4) Image's Past
One or two seconds before this picture was taken ___
• lady walks behind the bench
• they were standing ✔
• she got up on the table
• little one smiles pretty
5) Image's Future
One or two seconds after this picture was taken ___
• she put down the cat
• the bride dropped the bouquet
• she's gonna take a selfie
• the man closed the phone ✔
6) Person's Attribute
The person is ___
• dressed in same shirts
• a very young blonde boy ✔
• a girl in blue shorts
• barefoot and wearing a sundress
7) Person's Activity
The person is ___
• relaxing
• hosting a party
• carrying a chair ✔
• maintaining his skateboard
8) Person's Location
The person is ___
• in a house
• in a vehicle ✔
• in a bed
• in a room
9) Person – Object Relationship
The person is ___ the dog
• surfboarding with
• riding with the dog
• taking a photo with ✔
• playing with
10) Object's Attribute
The surfboard is ___
• yellow and flat ✔
• long and red
• blue and white
• white and gray
11) Object's Affordance
People could __ the bench
• place their luggage on
• sit on ✔
• take in the view
• relax on
12) Object's Position
The apple is __
• in a box
• on a stick ✔
• in a basket
• on the table

Fig. 8 Examples of answers selected using each individual cue as well as the full ensemble method. The first VGG column corresponds to the baseline feature (B.), while the following columns correspond to "B. + X" features following the same column order as in Table 1 (the Action feature is HICO).

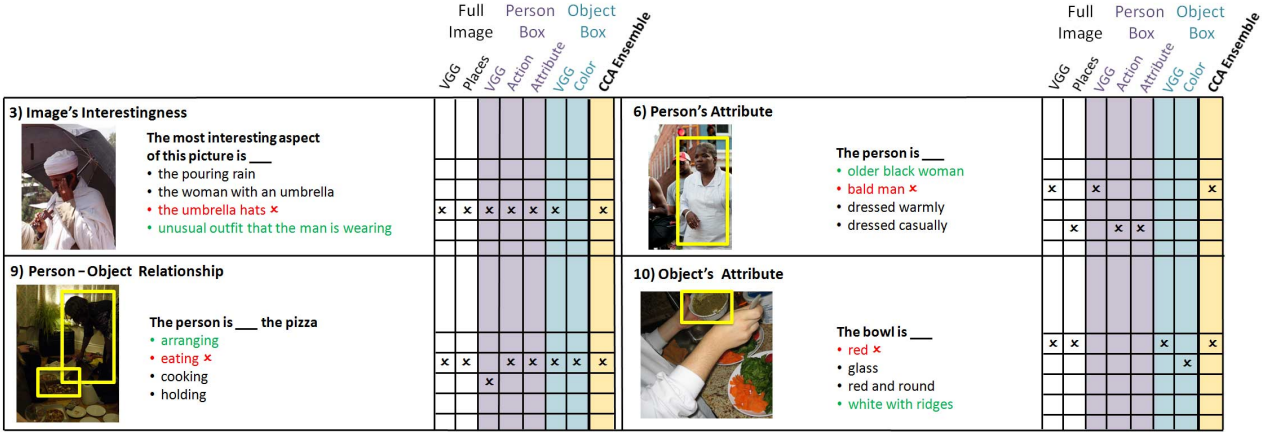

Fig. 9 Failure cases for four multi-choice question types from the Hard question-answering setting. Examples in the left column involve relatively rare concepts like "unusual outfit" and "arranging the pizza," while examples on the right are visually subtle or ambiguous.

Auer S, Bizer C, Kobilarov G, Lehmann J, Cyganiak R, Ives Z (2007) Dbpedia: A nucleus for a web of open data. In: ISWC/ASWC

Bourdev L, Maji S, Malik J (2011) Describing people: Poselet-based attribute classification. In: ICCV

Chao YW, Wang Z, He Y, Wang J, Deng J (2015) Hico: A benchmark for recognizing human-object interactions in images. In: ICCV

Duchi J, Shalev-Shwartz S, Singer Y, Chandra T (2008) Efficient projections onto the l1-ball for learning in high dimensions. In: ICML

Gao H, Mao J, Zhou J, Huang Z, Wang L, Xu W (2015) Are you talking to a machine? dataset and methods for multilingual image question answering. In: NIPS

Geman D, Geman S, Hallonquist N, Younes L (2015) Visual turing test for computer vision systems. PNAS 112(12):3618–23

Girshick R (2015) Fast r-cnn. In: ICCV

Gong Y, Ke Q, Isard M, Lazebnik S (2014) A multi-view embedding space for modeling internet images, tags, and their semantics. IJCV 106(2):210–233

Hardoon D, Szedmak S, Shawe-Taylor J (2004) Canonical correlation analysis; an overview with application to learning methods. Neural Computation 16(12):2639–2664

Hotelling H (1936) Relations between two sets of variables. Biometrika 28:312377

Jabri A, Joulin A, van der Maaten L (2016) Revisiting visual question answering baselines. In: ECCV

Lassila O, Swick RR (1999) Resource Description Framework (RDF) Model and Syntax Specification. Tech. rep., W3C, URL `http://www.w3.org/TR/1999/REC-rdf-syntax-19990222/`

Le QV, Mikolov T (2014) Distributed representations of sentences and documents. arXiv preprint abs/1405.4053

Lin TY, Maire M, Belongie S, Hays J, Perona P, Ramanan D, Dollár P, Zitnick CL (2014) Microsoft coco: Common objects in context. In: ECCV

Liu H, Singh P (2004) Conceptnet – a practical commonsense reasoning tool-kit. BT Technology Journal 22(4):211–226

Liu W, Anguelov D, Erhan D, Szegedy C, Reed S, Fu CY, Berg AC (2015) SSD: Single shot multibox detector. arXiv preprint abs/1512.02325

Lyu S (2005) Mercer kernels for object recognition with local features. In: CVPR

Malinowski M, Fritz M (2014) A multi-world approach to question answering about real-world scenes based on uncertain input. In: NIPS

Malinowski M, Rohrbach M, Fritz M (2015) Ask your neurons: A neural-based approach to answering questions about images. In: NIPS

Mallya A, Lazebnik S (2016) Learning models for actions and person-object interactions with transfer to question answering. In: ECCV

Mikolov T, Sutskever I, Chen K, Corrado GS, Dean J (2013) Distributed representations of words and phrases and their compositionality. In: NIPS

Mokarian A, Malinowski M, Fritz M (2016) Mean box pooling: A rich image representation and output embedding for the visual madlibs task. In: BMVC

Pishchulin L, Andriluka M, Schiele B (2014) Fine-grained activity recognition with holistic and pose based features. In: GCPR

Plummer BA, Wang L, Cervantes CM, Caicedo JC, Hockenmaier J, Lazebnik S (2015) Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. arXiv preprint abs/1505.04870

Ren M, Kiros R, Zemel R (2015a) Exploring models and data for image question answering. In: NIPS

Ren S, He K, Girshick R, Sun J (2015b) Faster r-cnn: Towards real-time object detection with region proposal networks. In: NIPS

Russakovsky O, Deng J, Su H, Krause J, Satheesh S, Ma S, Huang Z, Karpathy A, Khosla A, Bernstein M, Berg AC, Fei-Fei L (2015) ImageNet Large Scale Visual Recognition Challenge. IJCV 115(3):211–252

Shih KJ, Singh S, Hoiem D (2015) Where to look: Focus regions for visual question answering. In: CVPR

Shih KJ, Singh S, Hoiem D (2016) Learning to localize little landmarks. In: CVPR

Simonyan K, Zisserman A (2014) Very deep convolutional networks for large-scale image recognition. arXiv preprint abs/1409.1556

Socher R, Bauer J, Manning CD, Ng AY (2013) Parsing With Compositional Vector Grammars. In: ACL

Sudowe P, Spitzer H, Leibe B (2015) Person attribute recognition with a jointly-trained holistic cnn model. In: ICCV'15 ChaLearn Looking at People Workshop

Tandon N, de Melo G, Suchanek F, Weikum G (2014) Webchild: Harvesting and organizing commonsense knowledge from the web. In: ACM International Conference on Web Search and Data Mining

Tommasi T, Mallya A, Plummer B, Lazebnik S, Berg AC, Berg TL (2016) Solving visual madlibs with multiple cues. In: BMVC

Wang P, Wu Q, Shen C, van den Hengel A, Dick AR (2016) FVQA: fact-based visual question answering. arXiv preprint abs/1606.05433

Wu Q, Shen C, Hengel Avd, Wang P, Dick A (2016) Image captioning and visual question answering based on attributes and their related external knowledge. arXiv preprint abs/1603.02814

Xiao J, Hays J, Ehinger KA, Oliva A, Torralba A (2010) SUN database: Large-scale scene recognition from abbey to zoo. In: CVPR

Xu H, Saenko K (2015) Ask, attend and answer: Exploring question-guided spatial attention for visual question answering. arXiv preprint abs/1511.05234

Yang Z, He X, Gao J, Deng L, Smola AJ (2016) Stacked attention networks for image question answering. In: CVPR

Yu L, Park E, Berg AC, Berg TL (2015) Visual Madlibs: Fill in the blank Image Generation and Question Answering. In: ICCV

Zhou B, Lapedriza A, Xiao J, Torralba A, Oliva A (2014) Learning deep features for scene recognition using places database. In: NIPS

Zhou B, Tian Y, Sukhbaatar S, Szlam A, Fergus R (2015) Simple baseline for visual question answering. arXiv preprint abs/1512.02167

Zhu Y, Zhang C, Ré C, Fei-Fei L (2015) Building a large-scale multimodal knowledge base for visual question answering. arXiv preprint abs/1507.05670

Zhu Y, Groth O, Bernstein M, Fei-Fei L (2016) Visual7W: Grounded Question Answering in Images. In: CVPR